

RooStatsCms: a tool for analysis modelling, combination and statistical studies

D. Piparo¹, G. Schott¹ and G. Quast¹

¹ Faculty of Physics, Karlsruhe University, Wolfgang-Gaede-Str. 1, 76131 Karlsruhe

E-mail: danilo.piparo@cern.ch, gregory.schott@cern.ch, gunter.quast@cern.ch

Abstract.

RooStatsCms is an object oriented statistical framework based on the RooFit technology. Its scope is to allow the modelling, statistical analysis and combination of multiple search channels for new phenomena in High Energy Physics. It provides a variety of methods described in literature implemented as classes, whose design is oriented to the execution of multiple CPU intensive jobs on batch systems or on the Grid.

1. Introduction

Statistical analysis and the combination of measurements has a dominant importance in High Energy Physics. It is very challenging from the point of view of the tools to be deployed, the communication among the analysis groups and the definition of statistical guidelines. Previous experiments such as those at LEP [1] and Tevatron [2] already devoted huge efforts in this direction.

At the LHC, early results will require the combined analysis of different search channels and eventually the combination of results obtained by different experiments. There will definitely be a need to build complex models, i.e. parametrisations, to describe the experimental distributions, to quantify a possible signal excess in the data or to set a limit on the signal size in the absence of such an excess. In addition, a quantitative statistical treatment will require extensive studies based on Monte-Carlo, and should consider different statistical methods. The combination of analyses require a reliable and practical transfer of data and models across working group and experiment boundaries. Previous attempts to achieve these goals were built upon dedicated code for each analysis, and a very tedious and often error-prone transfer of the obtained results into the combination procedures.

In order to perform the statistical treatment of combination of analyses multiple methods are available. The choice of the method to use depends often on the context of the analysis and on the interpretation of the data by the experimenter. When multiple methods are applicable, a comparison of their results might be useful or is even required. It is therefore important to be able to easily switch between methods without too much effort. This was so far not possible as the implementation of different approaches were not unified in a single package and a comparison would require the user to learn how to use a number of packages.

It is the lack of a general, easy-to-use tool that led us to the decision to develop RooStatsCms (RSC) [3] for analysis modelling, statistical studies and combination of results. A first release of the RSC package has been provided in February 2008. It relies on the ROOT [4] extension

RooFit [5], from which it inherits the ability to efficiently describe the analysis model, thereby separating code and descriptive data, and easily perform toy Monte-Carlo experiments. A selection of different methods for statistical analysis and combination procedures is also included.

We start describing the software environment of RooStatsCms in section 2. Since RSC is made up of three components, i.e. the modelling part, devoted to the construction of the analysis model starting from a configuration file, the implementation of the statistical methods and a set of advanced graphics routines, a natural way to describe it is to characterise these parts separately, in sections 3, 4 and 5 respectively. In section 6, we then give an introduction to the statistical methods before showing some examples of applications of RSC in section 7.

2. Framework and software environment

RooStatsCms is entirely written in C++ and relies on ROOT. The ROOT Analysis Framework is the most widely used tool in the High Energy Physics community for data analysis. It exploits an advanced object oriented design to reach a high scalability and flexibility. Among its most important components there are cutting-edge visualisation tools, a rich set of container classes that are fully I/O aware, an extensive set of GUI classes, run-time object inspection capabilities, shared memory support and an automatic HTML documentation generation facility. The `TObject` class provides default behaviour and protocol for all objects in the ROOT system. It provides a protocol for object I/O, error handling, sorting, inspection, printing and drawing. Every object which inherits from `TObject` can be stored in the ROOT collection classes or written to ROOT files on disk. One of the key components of ROOT is the CINT C/C++ interpreter [6] since it allows rapid prototyping eliminating the typical time consuming edit/compile/link cycle. Existing C/C++ libraries can be easily interfaced to the interpreter. This is done by generating a dictionary from the functions and classes definitions which is then compiled and linked with the code into a single library. The CINT interpreter is fully embedded in ROOT allowing command line, scripting and programming languages to be identical. RSC is distributed with the CINT dictionaries and its classes and functions can therefore be used also inside macros or in the interpreter.

To reach a maximum flexibility and exploit all the recent technologies, we decided to use the RooFit toolkit. This package was a project started originally for the analyses of the BaBar experiment and is now a part of ROOT. The RooFit technology is based on a cutting-edge object oriented design, according to which almost every mathematical entity is represented by a class. For example, parameters and variables are treated symmetrically and can be expressed as `RooRealVar` objects, representing real intervals, holding an (asymmetric) error and a fit range. Also probability density functions (pdf) are represented by classes inheriting from `RooAbsPdf` and using their methods very complicated objects can be described. The numerous simple models provided in the package such as Gaussians, polynomials or Breit-Wigners can be combined to build the elaborate shapes needed for the analyses. Many operations from pdf objects are possible; the user can perform pdf addition, convolution or product of pdfs of different variables. The communication among the class instances bounded together in the complex structures that result from such operations, is granted by an advanced reference caching mechanism. The persistency of these composite objects is assured by the `RooWorkspace` class, which is able to go through these references and bring to disk all the necessary objects. In any case, RooFit takes care automatically of the normalisation of the pdfs with respect to all of their parameters and variables within their defined ranges. All the integration procedures are highly optimised, combining analytic treatment with advanced numerical techniques. Simultaneous and disjointed fits can be carried out with the possibility to interface RooFit to MINUIT and other minimisation packages. On the top of that, optimised toy Monte-Carlo dataset generation is provided.

RooStatsCms is integrated in the official CMS [7] Software Framework [8], in the `PhysicsTools/RooStatsCms` package, starting from the 3.1.X series.

3. Analyses modelling

In the analysis of a physics process the description of its signal and background components, together with correlations and constraints affecting the parameters, is a critical step. RooStatsCms provides the possibility to easily model the analyses and their combinations through the description of their signal and background shapes, yields, parameters, systematics and correlations via analysis configuration files, called *datacards*. The goal of the modelling component of RSC is to parse the datacard and generate from it a model according to the RooFit standards. There are a few classes devoted to this functionality, but the user really needs to deal only with one of them, the `RscCombinedModel`. The approach described above has mainly three advantages: the factorisation of the analysis description and statistical treatment in two well defined steps, a common base to describe the outcomes of the studies by the analysis groups, and a straightforward and documented sharing of the results.

A datacard is an ASCII file in the “.ini” format, therefore presenting key-value pairs organised in sections. This format was preferred to the XML because of its simplicity and high readability.

The parsing and processing of the datacard is achieved through an extension of the RooFit `RooStreamParser` utility class. This class is already rather advanced. Beyond reading strings and numeric parameters from configuration files, it implements the interpretation of conditional statements, file inclusions and comments. In presence of a complicated combination, the user can take advantage in RSC from these features specifying one single model per datacard file and then import all of them in a “combination card”.

Every analysis model can be described as a function of one or many observables, e.g. invariant mass, output of a neural network or topological information regarding the decay products. For each of these variables a description of the signal and background case is to be given, where both signal and background can be divided in multiple components, e.g. multiple background sources. To each signal and background components, a shape and a yield can be assigned. For what concerns the shape, a list of models is present and for those shapes which are not easily parametrisable, a TH1 histogram in a ROOT file can also be specified. The yields can be expressed as a product of an arbitrary number of single factors, for example: $\text{Yield} = \int \mathcal{L} \cdot \sigma \cdot BR \cdot \epsilon$, where $\int \mathcal{L}$ is the integrated luminosity, σ a production cross section, BR a decay branching ratio and ϵ is the detection efficiency.

Using RooFit, all the parameters present in the datacard can be specified as constants or defined in a certain range. In addition to that, exploiting the RSC implementation of the constraints, the user can directly specify the parameter affected by a Gaussian or a log-normal systematic uncertainties. In the former case, correlations can be specified among the parameters via the input of a correlation matrix.

In a combination some parameters might need to be the same throughout many analyses, e.g. the luminosity or a background rate. This feature is achieved in the modelling through a “same name, same pointer” mechanism. Indeed every parameter is represented in memory as a `RooRealVar` or, in presence of systematic uncertainties, as a derived object, the `Constraint` object and the `RscCombinedModel` merges all variables with the same name via an association to the same pointer.

4. Implementation of the statistical methods

Each statistical method in RooStatsCms is implemented in three classes types (this structure is reflected in the code by three abstract classes, see figure 1): the *statistical method* where all the time consuming operations such as Monte-Carlo toy experiments or fits are performed, the *statistical result* where the results of the computations are collected and the *statistical plot* whose role is to provide a graphical representation of the *statistical result*.

In many cases, e.g. frequentist approaches (see section 6.2), the CPU time needed for the calculations can be considerable. An interesting feature of the *statistical result* classes is that

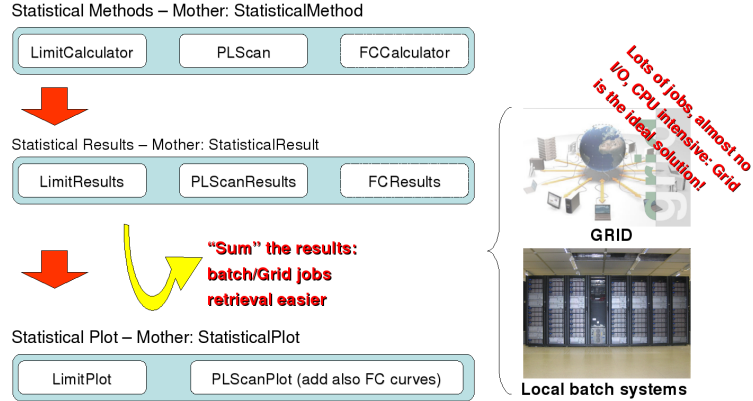


Figure 1. The RSC classes structure, excluding the modelling component. It is designed to ease job submission and recollection of results.

their objects can be “summed up”: this is very useful to accumulate statistics when combining the outputs of many processes. Indeed, the classes factorisation described above combined with the persistency of the RSC objects, eases the submission of jobs to a batch system or to the Grid and the recollection of the results, allowing to carry out such calculations at in reasonable timescales.

The *statistical plot* classes play a fundamental role in a statistical analysis, providing a graphical representation of the results in the form of self explanatory plots. The objects of these classes can directly be drawn onto a `TCanvas` via a `draw` method and, when needed, all the components used to produce the plot (`TGraph`, `TH1F`, `TLegend`, ...) can be saved separately in a ROOT file for a further manipulation.

5. Graphics routines

This category of classes is devoted to the production of plots. There are two types of plots covered: those that summarise information collected during the running of the statistical classes (such as figures 2 and 3), and plots allowing a graphical display of the physics results obtained. In this second category, if on the one hand, RooStatsCms does not provide any user-specific graphics routines, however, during the past decade, the LEP and Tevatron collaborations established a sort of standard to display the results of (combined) searches for new signals. This kind of plots are now well accepted in the community, and for this reason utility routines are provided to produce them. Figures 4 and 5 show two examples of plots. A few more examples are shown in section 7.

6. Statistical methods

We will now present the principles of some of the statistical methods implemented in the package before illustrating in the next section results of applications in CMS analyses. For a more detailed introduction to statistical methods in High Energy Physics see for example [9].

6.1. Profile likelihood approach

Suppose we have, for each of N events in a collection, a set of measured quantities $\underline{x} = (x_a, x_b, x_c, \dots)$ whose distributions are described by a joint probability density function, or pdf, $f(\underline{x}, \underline{\theta})$, where $\underline{\theta} = (\theta_1, \theta_2, \theta_3, \dots)$ is a set of K parameters. The likelihood function is then

defined by the equation:

$$L(\underline{x}, \underline{\theta}) = \prod_{i=1}^N f(x_i, \underline{\theta}). \quad (1)$$

To ease the calculations, the negative of the logarithm of the likelihood function $-\ln L$ (negative log-likelihood) is often used.

Focussing on a one-dimensional case, the profile likelihood method requires to perform a scan over a sensible range of values of the parameter of interest θ_0 . For every point of this scan, the value of θ_0 is fixed and $-\ln L(\theta_0)$ is minimised (i.e. a fit is performed) with respect to the remaining $K - 1$ parameters. The maximum likelihood estimator of the θ_0 parameter, noted $\hat{\theta}_0$, is the value where the negative log-likelihood function is at its minimum $-\ln L(\hat{\theta}_0)$.

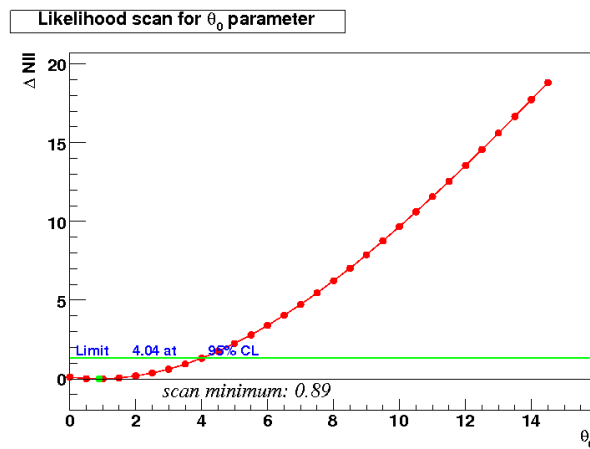


Figure 2. The negative log-likelihood scan over the parameter θ_0 . The horizontal line at $\Delta nll \simeq 1.36$ allows to read for the 95% CL upper limit on this parameter. The interpolation done by RooStatsCms between scan point pairs is linear, while the minimum of the scan is the minimum of a parabola built on the lowest three points.

Figure 2 shows an example of a profiled negative log-likelihood curve which has been offset by $-\ln L(\hat{\theta}_0)$. From this construction, it is possible to easily obtain the one- or two-sided confidence interval we are interested in. In the assumption of a parabolic shape of the negative log-likelihood function¹, the boundaries of a confidence intervals correspond to the values of θ_0 with

$$\Delta nll = -(\ln L(\theta_0) - \ln L(\hat{\theta}_0)) = \frac{n_\sigma^2}{2}, \text{ with } n_\sigma = \frac{\theta_0 - \hat{\theta}_0}{\sigma}. \quad (2)$$

where σ represent the Gaussian standard deviations. The mapping between the desired confidence level (CL) and the value of n_σ is given, in the Gaussian assumption, by the formulae:

$$n_\sigma = \sqrt{2} \cdot \text{Erf}^{-1}(CL) \quad (\text{two-sided}) \quad (3)$$

$$n_\sigma = \sqrt{2} \cdot \text{Erf}^{-1}(2 \cdot CL - 1) \quad (\text{one-sided}) \quad (4)$$

where Erf^{-1} is the Gaussian inverse error function. In the later case, an upper limit on the parameter θ_0 would then correspond then to the value $\theta_0 > \hat{\theta}_0$ obeying equations 2 and 4.

¹ It can be shown that this approach is also valid for a general scan shape [10].

6.2. Frequentist approach

Analysis of search results can be formulated in terms of hypothesis testing in a frequentist approach (for an explanation see [11]). We define H_b as the hypothesis that no signal is present over the background and H_{sb} the hypothesis that signal is also present. In order to quantify the degree in which each hypotheses are favoured or excluded by the experimental observation one chooses a test-statistics which ranks the possible experimental outcomes. A commonly used test statistics consist as the ratio of the likelihood function in both hypotheses: $Q = L_{sb}/L_b$ and the quantity $-2 \ln Q$ may also be used instead of Q . RooStatsCms also provides alternative choices for the test statistics such as the number of events or the profiled likelihood ratio.

A comparison of Q_{obs} for the data being tested to the probability distributions dP/dQ expected in both hypotheses allows to compute the confidence levels:

$$CL_{sb} = P_{sb}(Q < Q_{\text{obs}}), \quad \text{where} \quad P_{sb}(Q < Q_{\text{obs}}) = \int_{-\infty}^{Q_{\text{obs}}} \frac{dP_{sb}}{dQ} dQ, \quad (5)$$

$$CL_b = P_b(Q < Q_{\text{obs}}), \quad \text{where} \quad P_b(Q < Q_{\text{obs}}) = \int_{-\infty}^{Q_{\text{obs}}} \frac{dP_b}{dQ} dQ. \quad (6)$$

Small values of CL_{sb} (resp. CL_b) point out poor compatibility with the H_{sb} (resp. H_b) hypothesis and favour the H_b (resp. H_s) hypothesis. The functional form of the dP_{sb}/dQ and dP_b/dQ pdfs not being known a priori, a large amount of toy Monte-Carlo experiments are performed in order to determine it for two families of pseudo datasets: the ones in the signal+background and the ones in the background-only hypothesis (see figure 3).

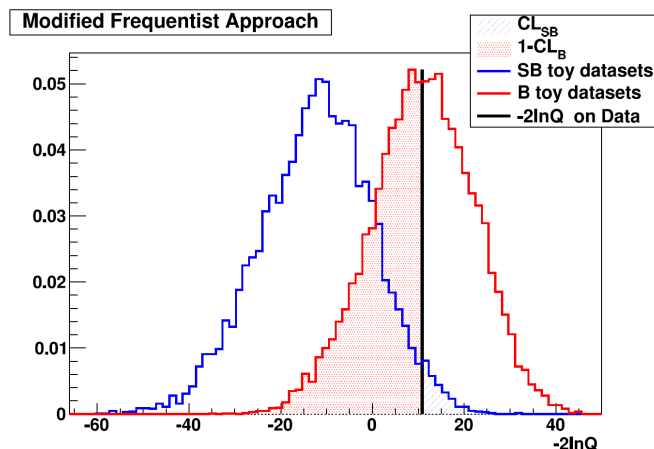


Figure 3. The distributions of $-2 \ln Q$ in the background-only (red, on the right) and signal+background (blue, on the left) hypotheses. The black line represents the value of $-2 \ln Q$ on the tested data. The shaded areas represent $1 - CL_b$ (red) and CL_{sb} (blue).

A significance estimation can be obtained using formula 4 on CL_{sb} . Moreover, the tested data can be said to be excluded at a given CL if $1 - CL_{sb}$ is smaller than this CL (or alternatively the CL_s prescription can be used (see below)). By varying the hypothesis being tested (for example varying the signal cross-section as on figures 4, 5, 6 and 7) one may also scan for the type of model that can be excluded with the given data. It should be observed that these confidence intervals do not have the same meaning as the ones obtained with the profile likelihood method or the Bayesian credibility intervals.

6.3. The CL_s prescription

Since the hypothesis being tested above is the signal *plus* background hypothesis and not the signal-only one, it is possible that unphysical regions are included in the confidence interval obtained, if the data is affected by large downward fluctuation of the background. In order to avoid this feature of a pure frequentist approach, the *modified*, or *conservative*, frequentist approach (CL_s method) was often used in High Energy Physics (e.g. by LEP, Tevatron and Hera experiments). Here, one uses the ratio of p-values, $CL_s = CL_{sb}/CL_b$, leading to more conservative limits. Even if CL_s is not technically a confidence level, the signal hypothesis is here considered excluded with a certain confidence level CL when $1 - CL_s < CL$.

6.4. Inclusion of systematics uncertainties

Systematics uncertainties can be taken into account by various techniques. In the likelihood methods described in section 6.1 a very convenient approach is to use the profiling procedure while in the frequentist method of section 6.2 a Monte-Carlo marginalisation technique can be applied. Both methods require to assume a probability distribution for the systematics, or *nuisance*, parameters (this probability distribution would be called the *prior probability* in a Bayesian context). RooStatsCms allows, for example, to assume a parameter θ_s to be distributed in an interval $[\theta_{s,\min}; \theta_{s,\max}]$ according to a Gaussian, a log-normal or a flat distribution.

In the same treatment, it is possible to take into account correlations between parameters by providing the full covariance matrix to RSC. While for uncorrelated nuisance parameters, the global prior probability distribution simply consist of the product of the individual distributions, when a correlation between $n \geq 2$ nuisance parameters is known, one can/should use a n -dimensional prior probability distribution. The profiling (see section 6.1) takes place through the minimisation of the negative log-likelihood function that could take into account the systematics and correlations. This does not require any Monte-Carlo integration. Suppose that the θ_s parameter is affected by the systematic uncertainties described by the $g(\theta_s)$ pdf. One writes the joint pdf describing the data and parameters as

$$f'(\underline{x}, \underline{\theta}) = f(\underline{x}, \underline{\theta}) \cdot g(\theta_s). \quad (7)$$

In the negative log-likelihood function $g(\theta_s)$ contributes as an additive penalty term: there is freedom in varying θ_s but $-\ln g(\theta_s)$ become large if going too far from the expected value (w.r.t. the magnitude of the assumed uncertainty). The scan of this altered likelihood preserves the position of the minimum point but has in general a larger curvature leading to broader confidence intervals and less aggressive limits. Once a dataset is specified, RooStatsCms, in presence of systematics, automatically creates a likelihood of the type described in equation 7.

The second approach applies Bayesian Monte-Carlo sampling described section 6.2 [12]. It consists in varying for each toy Monte-Carlo experiment the effective value of the nuisance parameters before generating the toy Monte-Carlo sample itself (that includes in addition the Poisson fluctuations). The whole phase space of the nuisance parameters is thus sampled through Monte-Carlo integration. The final net effect consist in a broadening of the $-\ln Q$ distribution and thus, as expected in presence of systematic uncertainties, a degraded separation of the hypotheses.

7. Examples of applications

RooStatsCms has been used in different contexts up to now and also within the CMS collaboration, exemplified here by some of the public results of Standard Model Higgs boson analyses [13, 14, 15]. Three examples of RSC applications are shown, two of which comprise a combination of analyses.

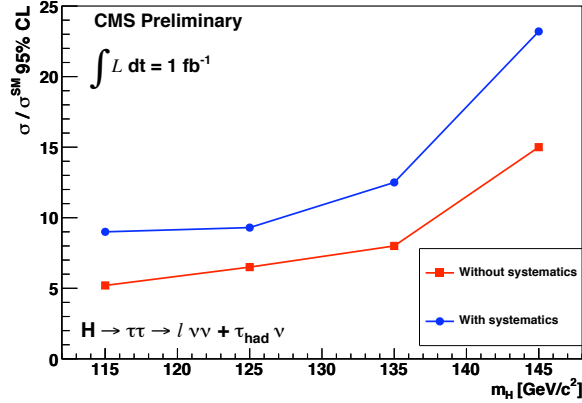


Figure 4. $H \rightarrow \tau\tau$ expected exclusion plot [13]: limits in presence or absence of systematics. The effect of the systematics is to deteriorate the exclusion power of the analysis.

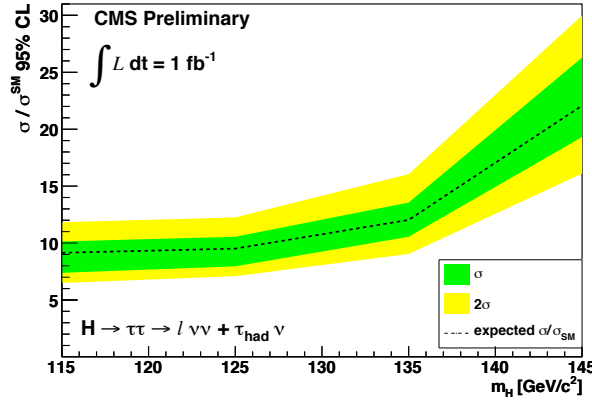


Figure 5. $H \rightarrow \tau\tau$ expected exclusion plot [13]: The 1σ and 2σ bands are obtained assuming a 1σ or 2σ upwards (downwards) fluctuation of the number of observed background events. The plot is produced with the `ExclusionBandPlot` class.

8. Conclusion and outlook

RooStatsCms is a tool for analysis modelling, statistical studies and the combination of analyses. It is based on ROOT and exploits the RooFit technology. The datacard approach described in section 3, provides a common base to describe the analyses and share the results among the groups. A number of popular statistical methods are provided by the package: frequentist and modified frequentist, profile likelihood, an interface for the BAT package providing Bayesian interpretation of the data [16] and a prototype for an implementation of the Feldman-Cousins method [17]. To perform the very lengthy calculations implied by some methods, the splitting of processes into sub-jobs and the recollection of results is eased due to the class structure. As shown in section 7, CMS has been using RSC in a number analyses. The implemented methods have been carefully validated by the analysis groups and analysis review committees, and the results are publicly available.

RooStatsCms can be also seen, in a wider context, as the starting point of the CMS

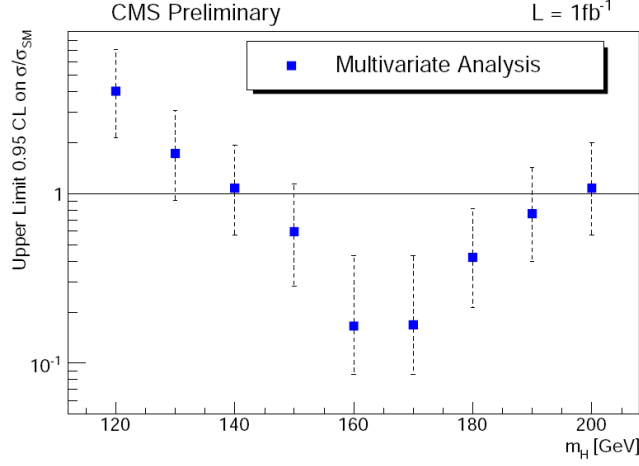


Figure 6. $H \rightarrow WW^{(*)} \rightarrow 2l2\nu$ [14]: expected 95% CL upper limits with the profile likelihood method for each of the Higgs mass hypotheses considered (in the assumption of no signal).

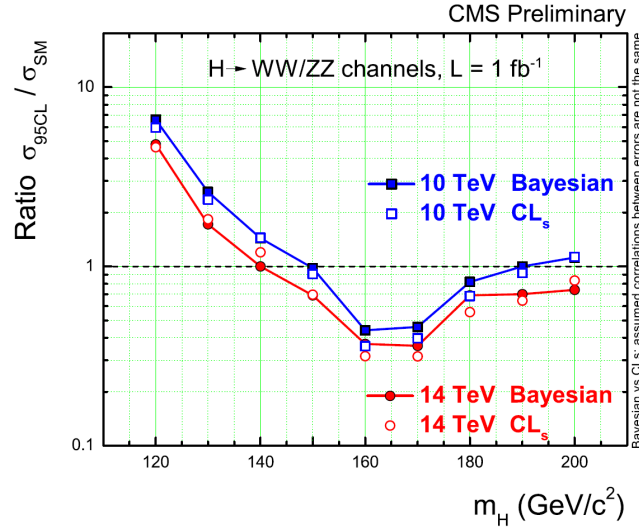


Figure 7. Projected exclusion limits for the Higgs boson at 14 and 10 TeV centre of mass energies: $H \rightarrow ZZ^{(*)}$ and $H \rightarrow WW^{(*)}$ channels [14, 15] are combined with two different methodologies, Bayesian and CL_s . Those results show a very reasonable agreement (even if the modelling of the constraints and correlations slightly differ in the two approaches). The Bayesian combination was carried out with a private tool, which confirmed the results obtained with RooStatsCms.

contribution to the RooStats [18] project: a joint effort of the LHC collaborations and the ROOT team, oversighted by a committee formed by the ATLAS [19] and CMS statistics committees. RooStats is part of ROOT since the 5.22 release of December 2008 and is currently in rapid evolution [20]. Parts of RSC are integrated into this first RooStats release, and further migration work is on-going. RSC presently implements more features than this initial RooStats version, but will be modified to use the future released versions of RooStats. The goal is to provide a common tool for statistical analysis and the combination of measurements implementing the methods

recommended by the statistics experts of the LHC collaborations. Furthermore, the tool will allow to store a full analysis model and transfer it from one working group to another or between experiments. This will ease enormously the cumbersome task of combination of experimental results. A full statistical re-analysis of the combined results, based on the original modelling of the contributing working groups, becomes possible while taking properly into account correlated parameters as well as common experimental or theoretical uncertainties. The anticipated broad basis of potential users will improve the reliability and robustness of the tool. Much experience towards a common statistics tool for High Energy Physics has been acquired throughout the process of the development of the CMS-specific tool RSC and the intense consultation with the experimental groups was extremely useful for the definition of the most important features and their implementation, thus set the basis for a decisive contribution of CMS to ROOT within the RooStats project. The CMS-specific tool will be adapted to rely on and interface the newly developed common classes, and will continue to provide a common interface for analysis modelling to the CMS collaboration. It will also be useful as a testing ground for new ideas to be explored within CMS.

Acknowledgements

We wish to thank the Higgs working group of CMS for their readiness to adopt the newly developed tool and for their valuable feedback. We are indebted to the statistics committee, and in particular to Bob Cousins for the exchange of ideas, advice and encouragement. We also wish to thank Wolfgang Wagner the possibility to validate some of the implemented methods by comparison to code used in one of the single top analyses by the CDF experiment. We thank the German Ministry of Science and Education, BMBF, for partly funding this work.

References

- [1] The LEP Electroweak Working Group, <http://lepewwg.web.cern.ch>.
- [2] The Tevatron Electroweak Working Group, <http://tevewwg.fnal.gov>.
- [3] D. Piparo, G. Schott, G. Quast, Proceedings of the 11th Topical Seminar on Innovative Particle and Radiation Detectors, Nuclear Physics B (Proceedings Supplement), arXiv:0812.2217.
See also: <http://www-ekp.physik.uni-karlsruhe.de/~RooStatsCms>.
- [4] R. Brun, F. Rademakers, "ROOT - An Object Oriented Data Analysis Framework", Nucl. Instrum. Meth. A389 (1997) 81. See also: <http://root.cern.ch>.
- [5] The RooFit Toolkit for Data Modeling, arXiv:physics/0306116, <http://roofit.sourceforge.net>.
- [6] M. Goto, "C++ Interpreter - CINT", CQ publishing, ISBN 4-789-3085-3 (Japanese).
- [7] The CMS collaboration, <http://cms.web.cern.ch>.
- [8] The CMSSW Application Framework,
<https://twiki.cern.ch/twiki/bin/view/CMS/WorkBookCMSSWFramework>.
- [9] F. James, "Statistical Methods in Experimental Physics 2nd Edition", Word Scientific 2006.
- [10] W.J. Metzger "Statistical Methods in Data Analysis", Katholieke Universiteit Nijmegen (2002).
- [11] A.L. Read "Modified Frequentist Analysis of Search Results (The CL_s Method)", CERN OPEN-2000-205.
- [12] R.D. Cousins and V.L. Highland, "Incorporating Systematic Uncertainties into an Upper Limit", Nucl. Instrum. Meth. A320 (1992) 331.
- [13] The CMS Collaboration, "Search for the Standard Model Higgs Boson Produced in Vector Boson Fusion and Decaying into a τ Pair in CMS with 1 fb⁻¹", CMS PAS HIG-08-008.
- [14] The CMS Collaboration, "Search Strategy for a Standard Model Higgs Boson Decaying to Two W Bosons in the Fully Leptonic Final State", CMS PAS HIG-08-006.
- [15] The CMS Collaboration, "SM Higgs to ZZ* to 4 Charged Leptons", CMS PAS HIG-08-003.
- [16] A.Caldwell, D. Kollar, K. Kröninger, "BAT - The Bayesian Analysis Toolkit", arXiv:0808.2552.
G. Schott, <http://www-ekp.physik.uni-karlsruhe.de/~schott/BCRooInterface>.
- [17] G.J. Feldman, R.D. Cousins "A Unified Approach to the Classical Statistical Analysis of Small Signals", Phys. Rev. D57 (1998) 3873.
- [18] The RooStats Project, <https://twiki.cern.ch/twiki/bin/view/RooStats>.
- [19] The ATLAS experiment, <http://atlas.ch>.
- [20] L. Moneta, "Overview of the New ROOT Statistical Software" (in these proceedings).